Introduction to Network Management

The field of network management encompasses a plethora of concepts, since it deals with all aspects of monitoring and fault-finding on the networks that underlie corporate systems. An interesting issue for any network manager is that it's very hard to draw the line between network management and the areas that surround it – server management, dealing with problems on the cabling infrastructure, application performance management and so on. For the moment, look at what one might call "pure" network management – proactively monitoring the corporate network

in order to identify issues as soon as they occur, or preferably to predict possible problems before they become reality by observing trends.

Management protocols

There are two standard protocols for network management and monitoring. The most important is the Simple Network Management Protocol (SNMP), which defines a set of instructions that can be used to view and change the settings on any SNMP-compliant device. The other is RMON, or the Remote Monitoring protocol, which bolts on to SNMP and allows remote collection of network information.

RMON

If you have a distributed network whose traffic levels you wish to monitor, you probably don't want to have a single central workstation interrogating all of the switches at all of your sites over WAN links – not only will it be hideously slow unless you have fast links, but it'll also have an impact on the loading of the WAN. The answer is to place an RMON "probe" (or more than one, if the network is extensive) at each of your sites, and allow it to collate the performance information of all of the SNMP devices on the LAN; instead of interrogating potentially dozens of switches and routers from the central office, you instead interrogate just your handful of RMON probes, thus reducing the impact on the network and the time you're spending waiting for data to download from remote sites.

SNMP

SNMP is a prehistoric protocol family, but its usefulness is that Noah ran it on his IBM abacus and so it's supported even by older hardware. Its drawbacks are that its security mechanisms suck and the data transfers you need to do to get stuff in and out are complex and long-winded, but this isn't a huge issue in these days of Gigabit networks and high-speed processors. SNMP is a standard that is continually added to as technologies develop. The accepted "standard" features are contained within the Management Information Base (MIB) definition, but part of this definition allows for manufacturers to define (and tag on) their own proprietary extensions to the standard MIB – which they do when, for instance, they invent a new technology which, by definition, isn't covered by the SNMP standard. As technologies become common, they are gradually absorbed into the standard MIB, and the evolution continues. There are three main facets to SNMP: viewing settings, changing settings and sending "traps". SNMP "get" The "get" commands of SNMP allow an application to pull information from a device. This could be the general configuration information (the link speed for each Token Ring port, for example, or the IP address of a router port) or it could be statistical information that accumulates during normal operation – packet counts, error counts and such like. SNMP "set" The "set" commands allow the network manager to write settings to a device – port speeds, IP addresses, VLAN definitions and so on. SNMP "trap" The majority of SNMP connections go from the management console (usually a Windows or X-Windows application of some kind) into a network device. The "trap" goes the opposite way, since it is SNMP's way of allowing a network device to alert a management station that there is a problem. So one might use a "set" command to tell a switch that you want to be alerted when you see more than 70% loading on an Ethernet segment for more than 10 seconds, and it'll use a "trap" to tell your management console in the event that this happens.

Management applications

SNMP in its own right has limited value – although you can use it to configure a device and receive performance and alerting information from it, the value is not in the raw data but what you can do with it. So to make the best use of SNMP, you'll need a network management program that collates all the information and presents you with a usable user interface that lets you do things like selecting multiple interfaces and applying a single operation to all items in the selection. Unfortunately, there are limitations to SNMP that restrict what you can do with a generic management program. For instance, although you could manage two different makes of 96-port switch through the same general application, the application probably couldn't figure out which ports are located on which card in the switch chassis – this is vendor-specific information that isn't available in SNMP. There are two options, then: either uses the vendor's own management package (which will be inherently aware of the structure of the chassis) or you can use a thirdparty package like HP OpenView which can accept vendor-specific plug-ins that fill the gap between the generic aspects of SNMP and the proprietary architecture of the devices you're managing. If you're fortunate enough to have the budget required to standardize on a single vendor, the vendor-specific package may well be the best choice, as it may well be customized toward proprietary management interfaces that the manufacturer has devised to sit alongside the SNMP ones; if you have a heterogeneous network, the generic-with-plug-in approach is the one to take.

Management capabilities

The extent to which you can monitor a network is defined entirely by the capabilities of the devices that drive that network. Although most routers include some kind of Simple Network Management Protocol (SNMP) support, it's common to find that companies have purchased the non-SMTP version of their switches, which makes it impossible to keep an eye on what volumes of traffic are traversing the LAN. It makes sense, then, to consider what management facilities you're likely to need when you buy your devices. As switches are a commodity these days, the cost is largely insignificant, and for medium- to high-end switches, the majority comes with SNMP thrown in any way as the extra cost to the vendor is minimal.

The physical layer

SNMP is mostly concerned with what's going on at ISO layers 2 and 3, but although the cabling infrastructure of the organization goes largely unchanged, it can develop faults. Manufacturers of high-end switches and routers such as Cisco have started to include cable-level diagnostics into their switches' port connectors, which can alert the network manager to wiring faults as they develop; if you're not a high-end installation, though, it makes sense to have some stand-alone cable diagnostic facilities to hand.

Extending above the network layers

Because the network's existence is due entirely to the fact that it's required to run the business' applications, there is an increasing trend to include an understanding of how the actual applications relate to the network. Application traffic analyzers, although traditionally regarded as a tool for developers of network applications, are becoming increasingly significant to the network manager because more and more business applications are becoming either Web Services or client-server applications – with an increased reliance on the network over the desktop applications of the past. By mapping the services an application uses onto the network, it is becoming increasingly possible to map from the network to the application and back again. For instance, if an application is experiencing a problem and that application is known to utilize a specific collection of servers, switches, routers and links, this makes the network manager's faultfinding task simpler. The reverse is also true: if the network manager observes an issue with a network device or segment, he understands what business applications this will influence.

Agents

The final thing we should touch on in an introduction to network management is the concept of an "agent" – a software application that resides on a device on the network and which is interrogated by a central management console from time to time. The SNMP modules in network devices are effectively "agents", since they reside in the devices and are contacted by the central console; for systems that don't support SNMP, you have the option of installing an "agent" on those systems. A typical example of this is server agents, who collate information such as CPU usage and network interface traffic – just like SNMP agents but generally using proprietary protocols to communicate with their central management devices.

Security management

Security management for networks is different for all kinds of situations. A home or small office may only require basic security while large businesses may require high-maintenance and advanced software and hardware to prevent malicious attacks from hacking and spamming.

Homes & Small Businesses

- A basic firewall or a unified threat management system.
- For Windows users, basic Antivirus software. An anti-spyware program would also be a good idea. There are many other types of antivirus or anti-spyware programs out there to be considered.
- When using a wireless connection, use a robust password. Also try to use the strongest security supported by your wireless devices, such as WPA2 with AES encryption.
- If using Wireless: Change the default SSID network name, also disable SSID Broadcast; as this function is unnecessary for home use. (However, many security experts consider this to be relatively useless).
- Enable MAC Address filtering to keep track of all home network MAC devices connecting to your router.
- Assign STATIC IP addresses to network devices.
- Disable ICMP ping on router.
- Review router or firewall logs to help identify abnormal network connections or traffic to the Internet.
- Use passwords for all accounts.
- For Windows users, Have multiple accounts per family member and use non-administrative accounts for day-to-day activities.
- Raise awareness about information security to children.

Medium businesses

- A fairly strong firewall or Unified Threat Management System
- Strong Antivirus software and Internet Security Software.
- For authentication, use strong passwords and change it on a bi-weekly/monthly basis.
- When using a wireless connection, use a robust password.
- Raise awareness about physical security to employees.
- Use an optional network analyzer or network monitor.
- An enlightened administrator or manager.

- Use a VPN, or Virtual Private Network, to communicate between a main office and satellite offices using the Internet as a connectivity medium. A VPN offers a solution to the expense of leasing a data line while providing a secure network for the offices to communicate. A VPN provides the business with a way to communicate between two in a way mimics a private leased line. Although the Internet is used, it is private because the link is encrypted and convenient to use. A medium sized business needing a secure way to connect several offices will find this a good choice.
- Clear employee guidelines should be implemented for using the Internet, including access to non-work related websites, sending and receiving information.
- Individual accounts to log on and access company intranet and Internet with monitoring for accountability.
- Have a back-up policy to recover data in the event of a hardware failure or a security breach that changes, damages or deletes data.
- Assign several employees to monitor a group like CERT which studies Internet security vulnerabilities and develops training to help improve security.

Large businesses

- A strong firewall and proxy to keep unwanted people out.
- A strong Antivirus software package and Internet Security Software package.
- For authentication, use strong passwords and change it on a weekly/bi-weekly basis.
- When using a wireless connection, use a robust password.
- Exercise physical security precautions to employees.
- Prepare a network analyzer or network monitor and use it when needed.
- Implement physical security management like closed circuit television for entry areas and restricted zones.
- Security fencing to mark the company's perimeter.
- Fire extinguishers for fire-sensitive areas like server rooms and security rooms.
- Security guards can help to maximize security.

School

- An adjustable firewall and proxy to allow authorized users access from the outside and inside.
- Strong Antivirus software and Internet Security Software packages.
- Wireless connections that lead to firewalls.
- Children's Internet Protection Act compliance. (Only schools in the USA)
- Supervision of network to guarantee updates and changes based on popular site usage.
- Constant supervision by teachers, librarians, and administrators to guarantee protection against attacks by both internet and sneakernetsources.
- An enforceable and easy to understand acceptable use policy which differentiates between school owned and personally owned devices
- FERPA compliance for institutes of higher education

Large government

- A strong firewall and proxy to keep unwanted people out.
- Strong antivirus software and Internet Security Software suites.
- Strong encryption.
- White list authorized wireless connection, block all else.
- All network hardware is in secure zones.
- All hosts should be on a private network that is invisible from the outside.
- Host web servers in a DMZ, or a firewall from the outside and from the inside.
- Security fencing to mark perimeter and set wireless range to this.
- Inventory controls of government owned mobile.

Types of Attacks

Networks are subject to attacks from malicious sources. Attacks can be from two categories "Passive" when a network intruder intercepts data traveling through the network, and "Active" in which an intruder initiates commands to disrupt the networks normal operation. Types of attacks include:

- Passive
- Network
- wiretapping
- Port scanner
- Idle scan
- Active
- Denial-of-service attack
- Spoofing
- Man in the middle
- ARP poisoning
- Smurf attack
- Buffer overflow
- Heap overflow
- Format string attack
- SQL injection

Firewall

A firewall is a device or set of devices designed to permit or deny network transmissions based upon a set of rules and is frequently used to protect networks from unauthorized access while permitting legitimate communications to pass.

Many personal computer operating systems include software-based firewalls to protect against threats from the public Internet. Many routers that pass data between networks contain firewall components and, conversely, many firewalls can perform basic routing functions.

Types

There are different types of firewalls depending on where the communication is taking place, where the communication is intercepted and the state that is being traced.

Network layer and packet filters

Network layer firewalls, also called packet filters, operate at a relatively low level of the TCP/IP protocol stack, not allowing packets to pass through the firewall unless they match the established rule set. The firewall administrator may define the rules; or default rules may apply. The term "packet filter" originated in the context of BSD operating systems.

Network layer firewalls generally fall into two sub-categories, stateful and stateless. Stateful firewalls maintain context about active sessions, and use that "state information" to speed packet processing. Any existing network connection can be described by several properties, including source and destination IP address, UDP or TCP ports, and the current stage of the connection's lifetime (including session initiation, handshaking, data transfer, or completion connection). If a packet does not match an existing connection, it will be evaluated according to the rule set for new connections. If a packet matches an existing connection based on comparison with the firewall's state table, it will be allowed to pass without further processing.

Stateless firewalls require less memory, and can be faster for simple filters that require less time to filter than to look up a session. They may also be necessary for filtering stateless network protocols that have no concept of a session. However, they cannot make more complex decisions based on what stage communications between hosts have reached.

Modern firewalls can filter traffic based on many packet attributes like source IP address, source port, destination IP address or port, destination service like WWW or FTP. They can filter based on protocols, TTL values, netblock of originator, of the source, and many other attributes.

Commonly used packet filters on various versions of Unix are *ipf* (various), *ipfw* (FreeBSD/Mac

OS X), *pf* (OpenBSD, and all other BSDs),*iptables*/*ipchains* (Linux).

Application-layer

Application-layer firewalls work on the application level of the TCP/IP stack (i.e., all browser traffic, or all telnet or ftp traffic), and may intercept all packets traveling to or from an application. They block other packets (usually dropping them without acknowledgment to the sender). In principle, application firewalls can prevent all unwanted outside traffic from reaching protected machines.

On inspecting all packets for improper content, firewalls can restrict or prevent outright the spread of networked computer worms and trojans. The additional inspection criteria can add extra latency to the forwarding of packets to their destination.

Application firewalls function by determining whether a process should accept any given connection. Application firewalls accomplish their function by hooking into socket calls to filter the connections between the application layer and the lower layers of the OSI model. Application firewalls that hook into socket calls are also referred to as socket filters. Application firewalls work much like a packet filter but application filters apply filtering rules (allow/block) on a per process basis instead of filtering connections on a per port basis. Generally, prompts are used to define rules for processes that have not yet received a connection. It is rare to find application firewalls not combined or used in conjunction with a packet filter.

Also, application firewalls further filter connections by examining the process ID of data packets against a ruleset for the local process involved in the data transmission. The extent of the filtering that occurs is defined by the provided ruleset. Given the variety of software that exists, application firewalls only have more complex rulesets for the standard services, such as sharing services. These per process rulesets have limited efficacy in filtering every possible association that may occur with other processes. Also, these per process ruleset cannot defend against modification of the process via exploitation, such as memory corruption exploits. Because of these limitations, application firewalls are beginning to be supplanted by a new generation of application firewalls that rely on mandatory access control (MAC), also referred to as sandboxing,
to protect vulnerable services. An example of a next generation application firewall is AppArmor included in some Linux distributions.

Proxies

A proxy device (running either on dedicated hardware or as software on a general-purpose machine) may act as a firewall by responding to input packets (connection requests, for example) in the manner of an application, while blocking other packets.

Proxies make tampering with an internal system from the external network more difficult and misuse of one internal system would not necessarily cause a security breach exploitable from outside the firewall (as long as the application proxy remains intact and properly configured). Conversely, intruders may hijack a publicly-reachable system and use it as a proxy for their own purposes; the proxy then masquerades as that system to other internal machines. While use of internal address spaces enhances security, crackers may still employ methods such as IP spoofing to attempt to pass packets to a target network.

Virtual LAN

A virtual local area network, virtual LAN or VLAN, is a group of hosts with a common set of requirements that communicate as if they were attached to the same broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical local area network (LAN), but it allows for end stations to be grouped together even if they are not located on the same network switch. VLAN membership can be configured through software instead of physically relocating devices or connections.

To physically replicate the functions of a VLAN, it would be necessary to install a separate, parallel collection of network cables and equipment which are kept separate from the primary network. However, unlike a physically separate network, VLANs must share bandwidth; two separate one-gigabit VLANs using a single one-gigabit interconnection can suffer both reduced throughput and congestion. It virtualizes VLAN behaviors (configuring switch ports, tagging frames when entering VLAN, lookup MAC table to switch/flood frames to trunk links, and untagging when exit from VLAN.)

Uses

VLANs are created to provide the segmentation services traditionally provided by routers in LAN configurations. VLANs address issues such as scalability, security, and network management. Routers in VLAN topologies provide broadcast filtering, security, address summarization, and traffic flow management. By definition, switches may not bridge IP traffic between VLANs as it would violate the integrity of the VLAN broadcast domain. This is also useful if someone wants to create multiple layer 3 networks on the same layer 2 switch. For example, if a DHCP server is plugged into a switch it will serve any host on that switch that is configured to get its IP from a DHCP server. By using VLANs you can easily split the network up so some hosts won't use that DHCP server and will obtain link-local addresses, or obtain an address from a different DHCP server.

VLANs are layer 2 constructs, compared with IP subnets which are layer 3 constructs. In an environment employing the VLANs, a one-to-one relationship often exists between VLANs and IP subnets, although it is possible to have multiple subnets on one VLAN. VLANs and IP subnets provide independent Layer 2 and Layer 3 constructs that map to one another and this correspondence is useful during the network design process.

By using VLANs, one can control traffic patterns and react quickly to relocations. VLANs provide the flexibility to adapt to changes in network requirements and allow for simplified administration.

Proxy server

In computer networks, a proxy server is a server (a computer system or an application) that acts as an intermediary for requests from clients seeking resources from other servers. A client connects to the proxy server, requesting some service, such as a file, connection, web page, or other resource available from a different server. The proxy server evaluates the request according to its filtering rules. For example, it may filter traffic by IP address or protocol. If the request is validated by the filter, the proxy provides the resource by connecting to the relevant server and requesting the service on behalf of the client. A proxy server may optionally alter the client's request or the server's response, and sometimes it may serve the request without contacting the specified server. In this case, it 'caches' responses from the remote server, and returns subsequent requests for the same content directly.

The proxy concept was invented in the early days of distributed systems as a way to simplify and control their complexity. Today, most proxies are a web proxy, allowing access to content on the World Wide Web.

Uses

A proxy server has a large variety of potential purposes, including:

- To keep machines behind it anonymous, mainly for security.
- To speed up access to resources (using caching). Web proxies are commonly used to cache web pages from a web server.
- To apply access policy to network services or content, e.g. to block undesired sites.
- To access sites prohibited or filtered by your ISP or institution.
- To log / audit usage, i.e. to provide company employee Internet usage reporting.

- To bypass security / parental controls.
- To circumvent Internet filtering to access content otherwise blocked by governments.
- To scan transmitted content for malware before delivery.
- To scan outbound content, e.g., for data loss prevention.
- To allow a web site to make web requests to externally hosted resources (e.g. images, music files, etc.) when cross-domain restrictions prohibit the web site from linking directly to the outside domains.

A proxy server that passes requests and responses unmodified is usually called a gateway or sometimes *tunneling proxy*.

A proxy server can be placed in the user's local computer or at various points between the user and the destination servers on the Internet.

A reverse proxy is (usually) an Internet-facing proxy used as a front-end to control and protect access to a server on a private network, commonly also performing tasks such as loadbalancing, authentication, decryption or caching.

Network operating system

A networking operating system (NOS), also referred to as the Dialoguer, is the software that runs on a server and enables the server to manage data, users, groups, security, applications, and other networking functions. The network operating system is designed to allow shared file and printer access among multiple computers in a network, typically a local area network (LAN), a private network or to other networks. The most popular network operating systems are Microsoft Windows Server 2003, Microsoft Windows Server 2008, UNIX, Linux, Mac OS X, and Novell NetWare.

Characteristics

Network Operating Systems are based on a client/server architecture in which a server enables multiple clients to share resources.

The Network Operating System can also do the following:

- Centrally manage network resources, such as programs, data and devices.
- Secure access to a network.
- Allow remote users to connect to a network.
- Allow users to connect to other networks like the Internet.
- Back up data and ensure its availability.
- Allow for simple additions of clients and resources.
- Monitor the status and functionality of network elements.
- Distribute programs and software updates to clients.
- Ensure efficient use of a server's capabilities.